

Control of critical behavior in a small-scale social system

Bryan C. Daniels,¹ David C. Krakauer,^{2,1} and Jessica C. Flack^{2,1}

¹*ASU-SFI Center for Biosocial Complex Systems,*

Arizona State University, Tempe, AZ

²*Santa Fe Institute, Santa Fe, NM*

Abstract

Over the last decade new technologies for making large numbers of fine-grained measurements have led to the surprising discovery that many biological systems sit near a critical point [1, 2]. These systems are potentially more adaptive in that small changes to component behavior can induce large-scale changes in aggregate structure and function. Examples include networks of neurons [3], ant groups cooperatively carrying a load [4], and animal groups forming flocks and schools [5]. Accounting for criticality remains a challenge as sensitivity to perturbation suggests a lack of robustness. Furthermore, change induced by perturbation may not be adaptive. Complicating matters further critical phenomena can result from history-dependent stochastic processes [6]. A question central to distinguishing among these conflicting views of criticality is to what degree criticality can be controlled by the components of the system [2]. We address the control of criticality using data on conflict dynamics and fight sizes from an animal society model system (*Macaca nemestrina*, $n=48$). The system is fundamentally finite so we operationalize criticality in information theoretic terms using Fisher information and a measure of instability. We analyze criticality using empirically-grounded equilibrium (maximum entropy) and dynamic (branching process) models of the monkeys' fight-joining behavior. We find that (1) this *heterogeneous, socially* organized system, like homogeneous, spatial systems (flocks and schools), sits near a critical point, (2) the contributions individuals make to how critical the system is can be quantified and vary, and (3) the distance from the critical point (DFC) can be controlled through biologically plausible mechanisms operating on this heterogeneity. These mechanisms include third-party policing, which dampens fight participation of the individuals with the largest effect on DFC [7–9]. Control of DFC allows biological systems to balance the tradeoff between robustness and need for rapid change.

We analyze a time-series of fights from a large, captive pigtailed macaque group collected over multiple observation periods during a four month period (SI Sec. I). The data consist of a series of binary fight participation vectors \vec{x} of length n . For each vector an individual is assigned a “1” if it participated in that fight and a “0” if it did not (SI, Sec. I).

To study criticality in our model system we can use tools from statistical mechanics. These tools are best deployed when the study system can be described within a simple, tractable, and well understood modeling framework. Hence our first task is to assess whether our data are consistent with any of three basic, but biologically valid, fight joining models: (1) independent fight joining decisions, (2) correlated fight joining decisions (equilibrium model), and (3) strategic, correlated fight joining decisions (dynamical branching process) (SI, Sec. II) [10]. We evaluate these models by determining how effectively each recovers a key social feature — the distribution of fight sizes s [11] (SI, Sec.II)—when parameterized by the empirical data.

In the independent model individuals join fights without taking into account which others are currently fighting (SI, Sec. IV).

The correlated decision-making model is an equilibrium maximum entropy model that fits all pairwise correlations and corresponds to a spin-glass model [12]. The resulting probability distribution over possible fights has relative negative log-likelihood,

$$L(\vec{x}) = - \sum x_i J_{ij} x_j, \quad (1)$$

where the coefficients J_{ij} are numerically fit to match individual and pairwise frequencies $\langle x_i \rangle$ and $\langle x_i x_j \rangle$.

In the dynamical branching process the dominant causes of conflict are temporal pairwise interactions—an individual joins the current fight with some finite probability only when it sees another individual join. Fight initiation is assumed to occur at a slower timescale, when a random individual becomes aggressive. Individuals join the fight by receiving aggression from or initiating aggression against an individual already in the fight. Parameters include individual initiation parameters p_{0i} , each denoting the relative probability that individual i begins a fight, and conditional redirection parameters p_{ij} , each denoting the probability that j joins a fight due to i having just joined. We consider only the relative ordering of individuals joining fights, not direct interactions, and estimate parameters by fitting conditional frequencies P_{ij} — the frequency of seeing individual j at any later time in a

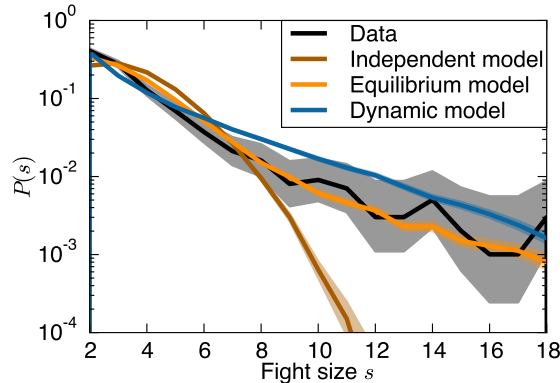


FIG. 1. Testing fight-joining processes fit to individual-level data: Models that include social correlations (the equilibrium maximum entropy and dynamic branching process models) can reproduce the relatively long tail of the observed distribution of fight sizes, whereas assuming independent fight joining events (the independent model) cannot.

fight sequence given that i appeared first.

We test the performance of each model by its ability to predict the distribution of fight sizes (Fig. 1) and correlations up to 3rd order (Fig. 7) (SI, Sec. IV). We rule out the independent model (SI, Sec. IV). We find the empirically parameterized maximum entropy (as in prior work [12]) and branching process models recover the observed distribution of fight sizes, indicating that these models are mechanistically consistent with the data. Because these models are well understood modeling frameworks in statistical physics, they offer a rigorous foundation for proceeding with a statistical mechanical description of the system. This allows us to define and quantify the sensitivity to perturbations and characterize criticality in a finite system.

A phase transition (SI, Sec V) can be thought of as sensitivity that diverges in the limit of infinite number of individuals due to a collective instability. This instability is an aggregate-level property that causes perturbations to individuals to be amplified and spread to change the behavior of the entire system (SI, Sec V C). While phase transitions are typically identified by examining the asymptotic behavior of the infinite limit, social systems are often not well approximated by such a “thermodynamic” limit. Yet sensitivity and instability are still important concepts in these systems. Here, in both the maximum entropy and dynamical branching models of our finite system, we find that changes to key parameters can lead quickly to increased sensitivity and instability. We take this to be an

operational definition of a “critical point” in a finite system (for a review of other definitions, see [2]).

We operationalize sensitivity as the derivative of average fight size with respect to an individual’s agitation, averaged over individuals. In the equilibrium model, the control parameter associated with the average fight size $\langle s \rangle$ is an external “field” h_{ext} , which we interpret as uniformly increasing each individual’s agitation, making it more likely to become involved in fights. Adding an external field to Eq. (1), we have $L(\vec{x}) = -\sum x_i J_{ij} x_j - h_{\text{ext}} \sum x_i$. The corresponding sensitivity is equal to the susceptibility χ :

$$\frac{1}{n} \sum_i \frac{d\langle s \rangle}{dh_i} = \chi = \frac{1}{n} \frac{d\langle s \rangle}{dh_{\text{ext}}}. \quad (2)$$

Results for the equilibrium model are shown in Fig. 2A. For fit parameters ($h_{\text{ext}} = 0$), there is increased sensitivity compared to a noninteracting model with the same mean fight size, meaning that an amplification process is occurring that makes changes to patterns of aggression at the individual level “visible” at the global system level. This amplification process cannot be attributed to external events as these data were collected in a captive setting in which such disturbances were minimized (SI, Sec II).

The susceptibility can also be interpreted as a Fisher information, an information theoretic quantity that describes how sensitive a distribution is to the parameters that describe it (SI VII). A large χ implies faster learning: large susceptibility means that aggregate level statistics are informative about conflict dynamics at the individual level [13, 14].

Analogously to the susceptibility in the equilibrium model, we can define sensitivity in the dynamic model as how quickly fight sizes grow following perturbation of redirection probabilities p_{ij} :

$$\frac{1}{n(n-1)} \sum_i \sum_{j \neq i} \frac{d\langle s \rangle}{dp_{ij}} \equiv \chi_{\text{dyn}} = \frac{1}{n(n-1)} \frac{d\langle s \rangle}{dp}, \quad (3)$$

where p adds probability uniformly to all redirection probabilities.

Fig. 2 suggests that the system is near a transition. To quantify this statement in a biologically meaningful way we measure how sensitivity and stability vary as we force—through parameterized-simulation—some number of individuals to fight continuously (while allowing the system to consist only of the non-forced individuals). The number of individuals required to reach a peak in sensitivity corresponding to a collective instability (SI, Sec.V C) is a measure of the *distance of the system from the critical point* (DFC). Beyond this peak lies saturation and decreasing sensitivity.

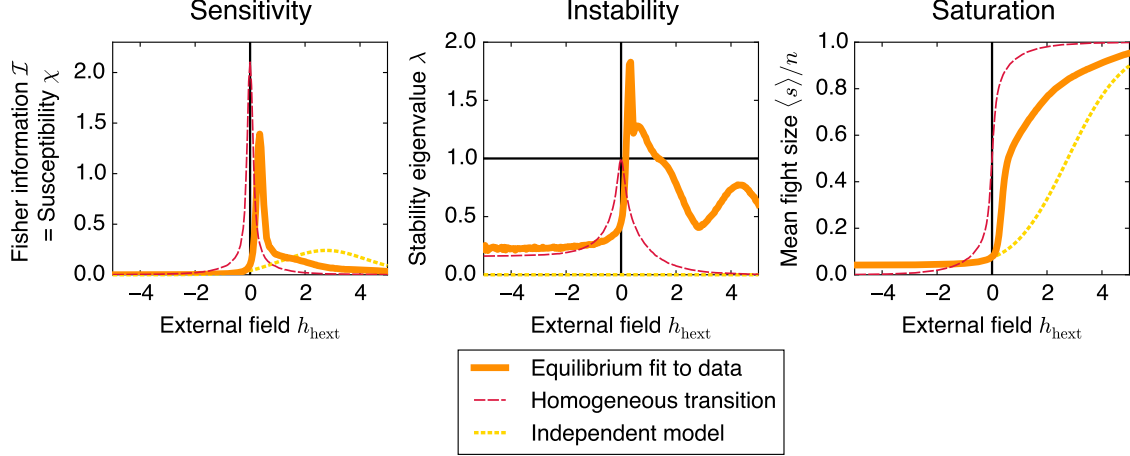


FIG. 2. The system is near a sensitive and unstable region of parameter space. (Left) The susceptibility χ , measuring the sensitivity of average fight size to an external perturbation, has a peak near the fit parameters (at $h_{\text{ext}} = 0$). The susceptibility is also equal to the Fisher information $F(h_{\text{ext}})$, describing how quickly the distribution over fights becomes distinguishable as h_{ext} is varied. (Middle) Instability of the lowest order mean-field solution is indicated by the eigenvalue λ becoming larger than 1. Yellow dashed lines indicate a system of the same size with independent individuals, and red dashed lines indicate a homogeneous system of the same size tuned to be marginally unstable at $h_{\text{ext}} = 0$. (Right) The transition is associated with large changes in mean fight size.

In the equilibrium model forcing is accomplished by removing the forced individuals and adding their interaction terms to the fields acting on the remaining individuals: the new fields are given by $J_{ii}^* = J_{ii} + \sum_{\text{forced } k} 2J_{ik}$. In the dynamic model this is accomplished by explicitly including the forced individuals at each timestep in the simulation but only recording the behavior of the remaining individuals. The change in sensitivity caused by forcing each individual also provides “sensitivity scores” that measure the extent to which each individual brings the system closer to criticality.

As shown in Fig. 3, both models predict a substantial increase in sensitivity when a few individuals are forced to be continuously involved in conflict. The system begins to saturate in sensitivity after 3-5 individuals are simultaneously forced. This result is similar to the finding that bird flocks require only a small proportion of informed individuals to correctly choose direction[15, 16].

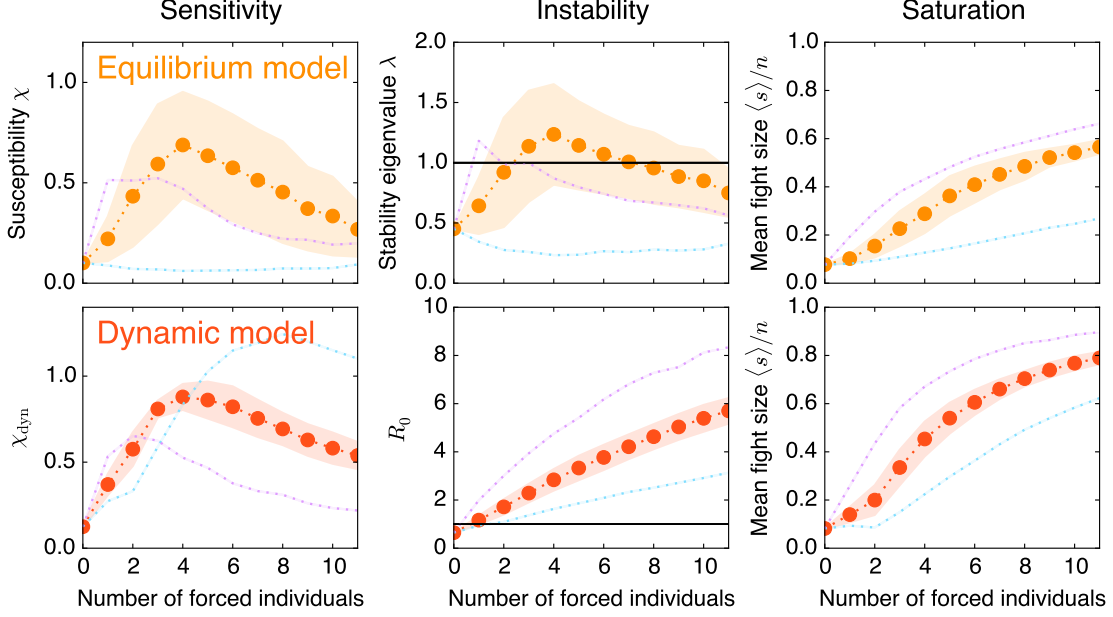


FIG. 3. In equilibrium (top) and dynamical (bottom) models, forcing a few individuals to become simultaneously aggressive leads to a more sensitive system. (Left) In each model, peak sensitivity is reached when on average 3-5 individuals are forced. Shaded regions indicate the standard deviation around the mean over different realizations of the chosen individuals. Purple and blue dotted lines indicate choices of forced individuals that maximize and minimize the resulting mean fight size. (Middle) The mean fight size. (Right) The lowest-order stability measured in each model, corresponding to the eigenvalue λ_0 in the equilibrium model and R_0 in the dynamical model (SI, Sec V C). A value above 1 in each case indicates instability to linear order.

We use perturbation theory to demonstrate that this sensitivity arises from instability in each model (SI, Sec.V D). The number of individuals required to reach a peak in sensitivity corresponding to an aggregate-level instability (SI, Sec.V C) is a *biologically meaningful measure* of the *distance of the system from the critical point*.

In addition we find that individuals vary in their sensitivity scores such that forcing participation of individuals with the top two to three sensitivity scores is sufficient to move the system close to a critical point (SI, Sec.VI).

These results tell us that fight involvement by the high sensitivity score individuals can be tuned two ways: through *direct tuning*, as they opportunistically or strategically increase or decrease their fight participation, and *indirect tuning*, through the actions of others that

increase or decrease the frequency with which aggression is directed or redirected at these individuals (SI, Sec.I and SI, Secs.VIA). This includes through the policing mechanism in which powerful third parties to the fight intervene and impartially break it up (SI, Sec.I and SI, Secs.VIA) [9]. Policing dampens fight involvement by significantly reducing how frequently fights erupt and also by reducing the frequency of redirected aggression received, particularly by high sensitivity individuals[7]. Hence DFC can in principle be controlled through policing or by the high sensitivity individuals themselves as they up or down regulate the frequency with which they join fights.

All biological systems need to balance stability and robustness with the need for rapid adaptive change. Yet many biological systems are observed to sit near a critical point, which would seem to suggest a lack of robustness. This apparent conflict can be resolved by the discovery that DFC can be tuned through realistic social mechanisms. Tuning DFC allows for switching between stability and criticality, providing a means for accessing alternative social structures that may be more appropriate if and when the environment should change[17]. This discovery raises many new questions: It is one thing for tuning to be possible and another to know when to tune and whether to tune towards robustness or criticality (SI, Sec.VIB). We propose robustness is a good strategy when the environment is stable and low variance. Criticality is a good strategy when the environment is uncertain (SI, Sec.VIB). In addition for tuning to be adaptive the state of the environment must be accurately perceived by, or mirrored in, the individuals whose behavior changes DFC. This appears to be the case in our system (SI, Sec.VIB) but may present a crucial evolutionary challenge for other systems. Future comparative studies are required to quantify the range of DFC across social groups within a species, as well as across biological systems more generally, and to study how DFC might be controlled in other systems. We can then assess whether variation in DFC and its control are correlated with rate of change in the environment and or environmental uncertainty as we expect.

END NOTES

- The authors thank Chris Ellison, Philip Poon, Eddie Lee, Eleanor Brush, Yoav Kallus, and Raissa D’Souza for helpful discussion. Partial support for this project was provided by the Templeton Foundation through grants to study complexity and the mind-brain

problem, and by the Office of Army Research under contract W911NF-13-1-0340. JCF thanks Frans de Waal and the animal care staff at Yerkes National Primate Research Center for support during data collection.

- BCD designed the study, built the models, analyzed and interpreted data, and wrote the paper. DCK contributed to model development, analysis, and interpretation, and wrote the paper. JCF designed the study, collected the data, contributed to model development, analysis, and interpretation, and wrote the paper. Correspondence and requests for materials should be addressed to Bryan Daniels (bryan.daniels.1@asu.edu).

-
- [1] Thierry Mora and William Bialek. Are Biological Systems Poised at Criticality? *Journal of Statistical Physics*, (May):268–302, June 2011.
- [2] Sergi Valverde, Sebastian Ohse, Malgorzata Turalska, Bruce J. West, and Jordi Garcia-Ojalvo. Structural determinants of criticality in biological networks. *Frontiers in Physiology*, 6:127, 2015.
- [3] David B. Kastner, Stephen A. Baccus, and Tatyana O. Sharpee. Critical and maximally informative encoding between neural populations in the retina. *Proceedings of the National Academy of Science*, 112:2533–2538, 2015.
- [4] Aviram Gelblum, Itai Pinkoviezky, Ehud Fonio, Abhijit Ghosh, Nir Gov, and Ofer Feinerman. Ant groups optimally amplify the effect of transiently informed individuals. *Nature Communications*, 6:7729, 2015.
- [5] William Bialek, Andrea Cavagna, Irene Giardina, Thierry Mora, Oliver Pohl, Edmondo Silvestri, Massimiliano Viale, and Aleksandra M. Walczak. Social interactions dominate speed control in poising natural flocks near criticality. *Proceedings of the National Academy of Sciences*, 111:7212–7217, 2014.
- [6] Bernat Corominas-Murtra, Rudolf Hanel, and Stefan Turner. Understanding zipf’s law with playing dice: history-dependent stochastic with collapsing sample-space have power law rank distributions. 2014.
- [7] Jessica C Flack, David C Krakauer, and Frans B M de Waal. Robustness mechanisms in

- primate societies: a perturbation study. *Proceedings of the Royal Society, Series B Biological sciences*, 272(1568):1091–1099, 2005.
- [8] Jessica C. Flack, M. Girvan, B.M. de Waal, Frans, and David C. Krakauer. Policing stabilizes construction of social niches in primates. *Nature*, 439:426–429, 2006.
 - [9] Jessica C. Flack, Frans B.M. de Waal, and David C. Krakauer. Social structure, robustness, and policing cost in a cognitively sophisticated species. *American Naturalist*, 165:E126–39, 2005.
 - [10] DC Krakauer, K Page, and JC Flack. The Immuno-Dynamics of Conflict Intervention in Social Systems. *PloS One*, 6:e22709, 2011.
 - [11] Simon DeDeo, David C. Krakauer, and Jessica C. Flack. Inductive game theory and the dynamics of animal conflict. *PLoS Computational Biology*, 6:e1000782, 2010.
 - [12] B. C. Daniels, D. C. Krakauer, and J. C. Flack. Sparse code of conflict in a primate society. *Proceedings of the National Academy of Sciences*, 109(35), August 2012.
 - [13] Martin Tchernookov and Ilya Nemenman. Predictive information in a nonequilibrium critical model. pages 1–16, 2012.
 - [14] Mikhail Prokopenko, Joseph T. Lizier, Oliver Obst, and X. Rosalind Wang. Relating Fisher information to order parameters. *Physical Review E*, 84(4):041116, October 2011.
 - [15] Iain D Couzin, Jens Krause, Nigel R Franks, and Simon a Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–6, feb 2005.
 - [16] N. E. Leonard, T. Shen, B. Nabet, L. Scardovi, I. D. Couzin, and S. a. Levin. Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences*, 109(1):227–232, 2012.
 - [17] Jessica C. Flack, Doug Erwin, Tanya Elliot, and David C. Krakauer. Timescales, symmetry, and uncertainty reduction in the origins of hierarchy in biological systems. *Cooperation and its Evolution*, pages 45–74, 2013.

Supplementary Information

Control of critical behavior in a small-scale social system

Bryan C. Daniels, David C. Krakauer, Jessica C. Flack

I. EMPIRICAL METHODS

A. Ethics statement

The data collection protocol was approved by the Emory University Institutional Animal Care and Use Committee and all data were collected in accordance with its guidelines for the ethical treatment of nonhuman study subjects.

B. Study system

The data were collected by JCF in 1998 from a large group of captive pigtailed macaques (*Macaca nemestrina*) socially-housed at the Yerkes National Primate Center in Lawrenceville, Georgia. Pigtailed macaques are indigenous to south East Asia and live in multi-male, multi-female societies characterized by female matrilineal and male group transfer upon onset of puberty [S1]. Pigtailed macaques breed all year. Females develop swellings when in Estrus. Macaque societies more generally are characterized by social learning at the individual level, social structures that arise from nonlinear processes and feed back to influence individual behavior, frequent non-kin interactions and multiplayer conflict interactions, the cost and benefits of which can be quantified at the individual and social network levels [S2–S7].

The study group contained $n = 48$ socially-mature individuals (we exclude non-mature individuals because their behavioral strategies are still developing and so are non-stationary over short timescales) and 84 individuals in total. Socially-mature males were at least 48 months and socially-mature females were at least 36 months by study start. These thresholds correspond to approximate onset of social maturity in pigtailed macaques. The study group had a demographic structure approximating wild populations and subadult and adult males were regularly removed to mimic emigration occurring in wild populations. All individuals, except 8 (4 males, 4 females), were either natal to the group or had been in the group since

formation. The group was housed in an indoor-outdoor facility, the outdoor compound of which was 125 x 65 ft.

Pigtailed macaques have frequent conflict and employ targeted intervention and repair strategies for managing conflict [S6]. Data on social dynamics and conflict were collected from this group over a stable, four month period. Operational definitions are provided below in SI Sec.I C.

C. Operational definitions

Fight: includes any interaction in which one individual threatens or aggresses a second individual. A conflict was considered terminated if no aggression or withdrawal response (fleeing, crouching, screaming, running away, submission signals) was exhibited by any of the conflict participants for *two minutes* from the last such event. A fight can involve multiple individuals. Third parties can become involved in pair-wise conflict through intervention or redirection, or when a family member of a conflict participant attacks a fourth-party. Fights in this data set ranged in size from 2 to 31 individuals, counting only the socially-mature animals. Fights can be represented as small networks that grow and shrink as pair-wise and triadic interactions become active or terminate until there are no more individuals fighting under the above described two minute criterion. In addition to aggressors, a conflict can include individuals who show no aggression or submission (*e.g.* third-parties who simply approach the conflict or show affiliative / submissive behavior upon approaching, and recipients of aggression who show no response to aggression (typically, threats) by another individual). Because conflicts involve multiple actors, two or more individuals can participate in the same conflict but not interact directly.

In this study only information about fight composition (which individuals were involved) is used. Only fights that included two or more socially-mature individuals were used in the analysis; the data includes $N = 994$ such fights. We do not consider internal aspects of the fight, such as who does what to whom, except for the order of each individual's first involvement in the fight (used to estimate time-ordered conditional probabilities for use in the dynamical branching process model). Time data were collected on fight onset and termination but are not used in these analyses.

Power: The degree of consensus among individuals in the group about whether an in-

dividual is capable of using force successfully [S3, S8]. In previous works we showed that consensus can be quantified by taking into account the total number of subordination signals an individual receives and multiplying this quantity by a measure of the diversity of signals received from its population of signalers (quantified by computing the Shannon entropy of the vector of signals received by individual i) [S3]. In pigtailed macaque societies, the subordination signal is the silent bared teeth display [S2] emitted outside the conflict context during pass-byes and affiliative interactions. The distribution of power in our study group is heavy tailed, such that a few individuals are disproportionately powerful.

Policing: A policing intervention is an impartial intervention performed by a third party into an ongoing conflict [S6]. Three males and one female perform the majority of effective policing interventions but only the three males (Eo, Qs, Fo) specialize on policing [S9]. These four individuals occupy the top four spots in the power distribution [S6, S8].

Redirection: A redirection occurs when an aggressor, recipient, or intervener directs aggression at a third (or fourth) individual who was not its original target or attacker. The target of the redirection may not have been involved in the fight until the redirection, or may have been involved in the fight but interacting with individuals other than the redirecting individual.

D. Data collection protocol

The data were collected by a trained observer (J.C. Flack). The observer spent roughly 100 hours prior to data collection learning to recognize individuals and accurately code their behavior from the observation tower above the monkey compound. Accuracy was validated by a second trained observer (F.B.M. de Waal). JCF also evaluated coding accuracy using video. Coding accuracy was nearly 100%.

During observations all individuals were confined to the outdoor portion of the compound and were visible to the observer. The ≈ 150 hours of observations occurred for up to eight hours daily between 1,100 and 2,000 hours over a twenty-week period from June through October 1998, and were evenly distributed over the day. Conflict and signaling data were collected using all-occurrence sampling in which the entire conflict event is followed from start to finish and all participants and their behavior are recorded.

Provisioning occurred before observations and once during observations at approximately

the same time each day. The group was stable during the data collection period (defined as no reversals in status signaling interactions resulting in a change to an individual's power score; see Ref. [S3]).

II. NOTES ON MODEL DESCRIPTIONS AND JUSTIFICATION

We first evaluate which of three basic, empirically-grounded fight joining models explain our macroscopic observable, the distribution of fight sizes. We only accept a model if it is both consistent with the measured, microscopic data and can recover in simulation the observed, measured macroscopic output. Hence all of our models are closely tethered to the measured data and biological details of our model system.

It is important to realize that the parameters in the maximum entropy and branching models come from the microscopic data. The models do not assume prescribed values for parameters but are perhaps better viewed as hypotheses about the ways in which the measured, microscopic detail is connected to observed macroscopic patterns.

All models assume that events external to the system do not create correlations in behavior. The data were collected in a controlled, captive setting designed to minimize the influence of external events, and we have no evidence for important, consistent external forcers of conflict.

A. Independent model description

The independent model assumes individuals do not respond to each other but instead join fights without regard to who else is fighting. In this case, perturbations to individual conflict behavior would have no additional effect on group behavior.

The independent model fails to recover both the observed distribution of fight sizes (Fig. 1) and the observed significant pairwise correlations (Fig. 7).

B. Pairwise maximum entropy model description

Individuals sometimes randomly join fights and sometimes the decision to fight reflects strategic interactions at a pairwise level[S10]. We capture this interpretation of the micro-

scopic data using a maximum entropy approach, which corresponds to the spin-glass model of magnetic systems in physics. Because the model is parameterized by the data it is empirically grounded and serves as a valid biological hypothesis. However we note that it is less mechanistically specific than the branching process described below: spin-glass interactions are not directional or time-ordered, but rather operate symmetrically and over the timescale of an entire fight bout.

The pairwise maximum entropy model, with parameters fit from the microscopic data, recovers well the distribution of fight sizes (Fig. 1). The good performance of the model leads to the prediction that the sensitivity χ is about twice that of a system with the same conflict frequencies but no strategic interactions.

C. Branching process description

Another reasonable interpretation of the microscopic data is that the random component of fight joining decisions is very small and the decision to fight reflects strategic interactions at a pairwise level. This leads to a branching process model that, with parameters fit from the microscopic data, also recovers the observed distribution of fight sizes. The collective behavior produced by this model can be simply understood in terms of a single parameter, the branching ratio. Additionally, branching process models like this one have been used in previous work to explore other aspects of conflict dynamics in this system, including the role of policing and other forms of third-party intervention in the infectivity of aggression [S9].

III. NOTES ON MODEL INFERENCE

A. Independent model inference

The independent model consists of individuals participating in conflict randomly, with the frequencies of individual appearance equal to their empirically measured, heterogeneous values $f_i = \langle x_i \rangle$. Naively, this can be written as a relative negative log-likelihood¹ $L^{\text{ind}}(x) =$

¹ The relative negative log-likelihood $L(x)$ of state x is related to the likelihood $p(x)$ by $p(x) = \exp(-L(x))/Z$, where Z is a normalization constant set by the constraint that the sum of likelihoods over all states is one. In statistical physics, Z is the partition function and $L(x)$ is proportional to the free energy of state x .

$\sum_i -h_i x_i$, and is equivalent to a maximum entropy model that matches only the frequencies f_i .

However, as detailed in SI Sec I C, a fight was operationalized for these analyses as involving two or more socially-mature individuals. (Observed fights that involved only juveniles and 0 or 1 mature individuals are therefore excluded.) Correspondingly, we forbid our models from producing fights of size smaller than two. We treat this as an additional constraint on the model. The resulting maximum entropy model is then the one in which the likelihood of states with fewer than two individuals present is taken to zero. This corresponds to a relative negative log-likelihood

$$L_{\alpha \rightarrow \infty}^{\text{ind}} = \lim_{\alpha \rightarrow \infty} \left[\sum_i h_i x_i + \alpha \Theta \left(2 - \sum_j x_j \right) \right], \quad (4)$$

where $\Theta(z)$ is 0 when $z \leq 0$ and 1 when $z > 0$.

In the unconstrained case ($\alpha = 0$), we can easily solve for h_i :

$$L_{\alpha=0}^{\text{ind}} = \sum_i h_i x_i \quad (5)$$

$$h_i = -\log \left(\frac{\langle x_i \rangle_{\alpha=0}}{1 - \langle x_i \rangle_{\alpha=0}} \right). \quad (6)$$

We must now solve numerically for h_i to match the empirically measured $f_i = \langle x_i \rangle = \langle x_i \rangle_{\alpha \rightarrow \infty}$. To accomplish this, note that the unconstrained model will have modified statistics:

$$\langle x_i \rangle_{\alpha=0} = (1 - f_0 - f_1) \langle x_i \rangle + f_{i1} \quad (7)$$

$$\langle x_i x_j \rangle_{\alpha=0} = (1 - f_0 - f_1) \langle x_i x_j \rangle,$$

where f_0 is the frequency of size zero fights, f_{i1} is the frequency of size one fights consisting solely of individual i , and $f_1 = \sum f_{i1}$ is the overall frequency of size one fights (all measured in the unconstrained model). In terms of unconstrained individual frequencies, these are

$$f_0 = \prod_i (1 - \langle x_i \rangle_{\alpha=0}) \quad (8)$$

$$f_{i1} = \langle x_i \rangle_{\alpha=0} \frac{f_0}{1 - \langle x_i \rangle_{\alpha=0}}. \quad (9)$$

We use an iterative procedure to solve equations (7) for $\langle x_i \rangle_{\alpha=0}$, which are then used in Eq. (6) to find the fields. Finally, samples from the independent model [Eq. (4)] are produced by sampling using Eq. (5) and simply discarding samples in which fewer than 2 individuals appear. For our data, this results in discarding about 17% of samples produced with Eq. (5).

B. Pairwise maximum entropy model inference

We next constrain our equilibrium maximum entropy model to match the frequencies of appearance of both individuals and pairs of individuals. This model is known to be the spin-glass Ising model [S11], with relative negative log-likelihood

$$L_{\alpha=0}^{\text{SG}}(x) = \sum_{ij} -x_i J_{ij} x_j. \quad (10)$$

As in the independent model, the likelihood of fights with fewer than two individuals is taken to zero:

$$L_{\alpha \rightarrow \infty}^{\text{SG}}(x) = \lim_{\alpha \rightarrow \infty} \sum_{ij} -x_i J_{ij} x_j + \alpha \Theta \left(2 - \sum_k x_k \right). \quad (11)$$

The statistics we fit in the equilibrium model are the individual and pairwise frequencies of appearance:

$$f_i = \langle x_i \rangle = \frac{N(x_i)}{N}, \quad (12)$$

$$f_{ij} = \langle x_i x_j \rangle = \frac{N(x_i, x_j)}{N} \quad \text{for } i \neq j, \quad (13)$$

with $N(x_i)$ and $N(x_i, x_j)$ representing, respectively, the observed number of appearances in unique fights of the individual i and the pair i, j .

We would like to fit the individual and pairwise frequencies of appearance to the precision that the data supports. As a measure of the goodness of fit, we use the average of normalized residuals

$$\langle \chi_{\text{SG}}^2 \rangle = \frac{2}{n(n+1)} \sum_i \sum_{j \geq i} \frac{(f_{ij}^{\text{data}} - f_{ij}^{\text{model}})^2}{\sigma_{f_{ij}}^2}, \quad (14)$$

where

$$\sigma_{f_{ij}}^2 = \frac{f_{ij}(1 - f_{ij})}{N} \quad (15)$$

is the expected variance of each residual due to finite N , and repeated indices are understood as individual frequencies: $f_{ii} = f_i$. A good fit is expected to have $\langle \chi^2 \rangle \approx 1$.

To perform fitting, we use a simplified method that starts with the mean field solution and varies a single parameter corresponding to weighting a non-interacting prior. Specifically, we make use of the L_2 -regularized mean field entropy of Ref. [S12]. The regularization consists of a Gaussian prior with a form designed to make the mean-field case easily solvable,

corresponding to an additional term in the relative negative log likelihood

$$L_{\text{prior}} = \gamma \sum_i \sum_{j>i} J_{ij}^2 f_i(1-f_i)f_j(1-f_j), \quad (16)$$

where γ is the strength of the prior, which favors interactions J_{ij} that are smaller in magnitude. The mean field solution under this regularization is [S12]

$$\begin{aligned} J_{ij}^{\text{MF}} &= \frac{J'_{ij}}{\sqrt{f_i(1-f_i)f_j(1-f_j)}}, \quad i \neq j \\ J_{ii}^{\text{MF}} &= \sum_{j \neq i} J_{ij}^{\text{MF}} \left((f_{ij} - f_i f_j) \frac{f_i - \frac{1}{2}}{f_i(1-f_i)} - f_j \right), \end{aligned} \quad (17)$$

where J' is the matrix that has the same eigenvectors v_q as the correlation matrix

$$C = \frac{f_{ij} - f_i f_j}{\sqrt{f_i(1-f_i)f_j(1-f_j)}} \quad (18)$$

and eigenvalues $j_q = 1/\hat{c}_q$, where \hat{c}_q are regularized versions of C 's eigenvalues c_q :

$$\hat{c}_q = \frac{1}{2} \left(c_q - \gamma + \sqrt{(c_q - \gamma)^2 + 4\gamma} \right). \quad (19)$$

This regularization is typically used in a Bayesian sense to avoid overfitting, where a typical value of the regularization strength is $\gamma = 1/(10Nf^2(1-f)^2)$, with N the number of samples and $f = n^{-1} \sum_i f_i$ the average individual frequency [S12]. Avoiding overfitting, however, still typically requires either enormous N (e.g. [S11]) or restricting the effective number of fit parameters via an expansion (e.g. [S12]). In our case, N is fundamentally limited in that we are describing a stable social epoch of finite duration. In addition, typical high-temperature expansions cannot easily incorporate the restriction that fights have a minimum number of participants [the α term in Eq. (11)].

Alternatively one can treat γ as a fitting parameter that interpolates between the mean-field solution (which we find overestimates the strength of interactions) and the case of independent individuals. Although it is not *a priori* obvious that varying this single parameter will be enough to fit the observed statistics within expected statistical fluctuations, we find that this is true for our data. Numerically sampling from the distribution defined by Eq. (11) with the regularized mean-field J from Eq. (17), we minimize $\langle \chi^2 \rangle$ from Eq. (14) as a function of γ . Sampling is performed using a standard Metropolis Markov Chain Monte Carlo approach. In evaluating the fit, we choose the number of samples to scale with the number of data samples, using $N_{\text{samples}} = 20N$.

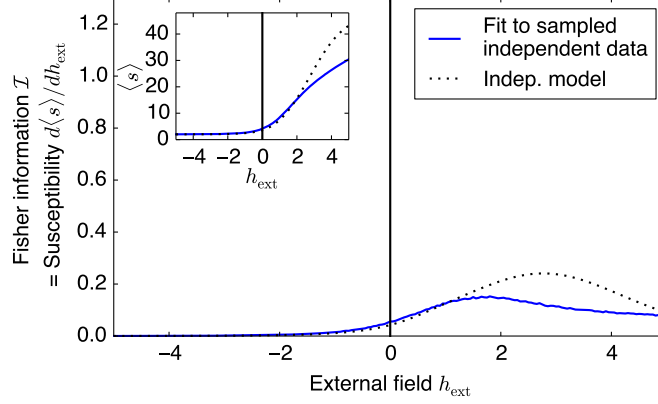


FIG. 4. Performing the second-order maximum entropy fitting procedure on an amount of data equal to the observed fights drawn from a noninteracting model produces a flat susceptibility curve comparable to the known exact susceptibility for this case.

As a simple check that this inference approach is not biased to find more sensitive systems, we infer a pairwise model using data produced by the independent model and compare its susceptibility to the known exact value [Eq. (30)] in Fig. 4. The resulting inferred model has susceptibility that stays close to the true value for small h_{ext} and remains smaller than the true value for larger h_{ext} .

C. Branching process inference

In the branching process model, we use time-ordered appearance data, fitting the time-ordered conditional appearance probabilities

$$P_{ij} = \frac{N(x_i^t, x_j^T)}{N(x_i)}, \quad (20)$$

where $N(x_i^t, x_j^T)$ counts the number of times individual j appears in the same fight bout as individual i , but at a later time $T > t$.²

The parameters we vary in the heterogeneous branching process model are the single-step conditional probabilities p_{ij} , which measure the probability that individual j appears in step $t+1$ of the branching process given that individual i appeared in step t . Being probabilities,

² Note that this is different than previous work in Inductive Game Theory [S13]: there (time directed) correlations were measured between individual appearances in separate fight bouts, whereas here we measure correlations within fight bouts.

p_{ij} are constrained to values $0 \leq p_{ij} \leq 1$. (Note that the branching model is thus limited in the extent to which it can represent inhibitory interactions, *e.g.* if i 's involvement in the fight deters j from joining.) We modify this constrained optimization problem into an unconstrained one by defining $\bar{p}_{ij} = |\tanh^{-1} p_{ij}|$ and performing the optimization over the (unconstrained) \bar{p}_{ij} parameters.

In the branching process simulation, the first individual to join each fight bout is chosen randomly with probability proportional to the frequency with which each individual appears at the beginning of fights in the data. At each subsequent time step in the branching process, each individual j who has not yet been activated in the current fight has a probability of joining equal to the sum of p_{ij} for all i active in the previous time step. The fight bout concludes when no individuals are active in a given time step. As discussed above, fight bouts that do not grow beyond a single individual are discarded.

Analogously to the case of the equilibrium maximum entropy model, the branching process parameters are fit by minimizing

$$\langle \chi_{\text{Branching}}^2 \rangle = \frac{1}{n(n-1)} \sum_i \sum_{j \neq i} \frac{(P_{ij}^{\text{data}} - P_{ij}^{\text{model}})^2}{\sigma_{P_{ij}}^2}, \quad (21)$$

where

$$\sigma_{P_{ij}}^2 = \frac{P_{ij}(1 - P_{ij})}{N(x_i)}. \quad (22)$$

We use a standard Levenberg-Marquardt algorithm (`scipy.optimize.leastsq`) that uses individual residuals and a lowest-order approximation for the Jacobian with respect to parameters. We find that restarting the Levenberg-Marquardt routine every 10 steps (effectively resetting its damping parameter to avoid unnecessarily small steps arising from its assumption of a non-stochastic objective function) produces faster convergence with fewer samples required for each estimate of the residuals and Jacobian. Minimization is stopped once $\langle \chi^2 \rangle$ defined in Eq. (21) falls below 1. In addition, we find that switching between simple gradient descent steps and Levenberg-Marquardt minimizations allows for more efficient fitting of larger P_{ij} .

The resulting inferred redirection probabilities p_{ij} are visualized in Fig. 5.

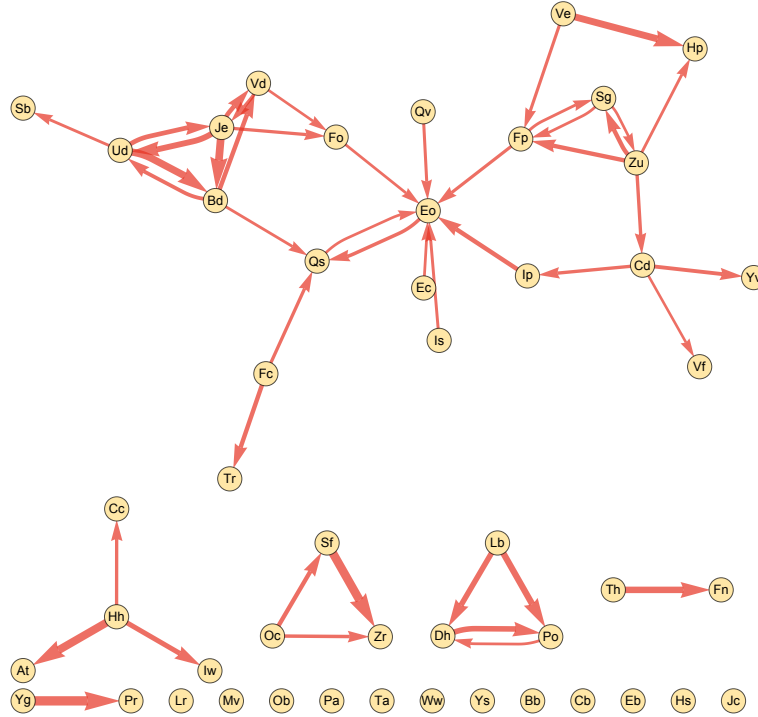


FIG. 5. A visualization of the inferred branching process graph. In the full branching process model, all possible directed pairwise interactions are included, but here we visualize the most important interactions by displaying those with triggering probabilities > 0.09 . The thickness of each arrow corresponds to the probability of triggering, with the thickest lines correspond to a probability of about $1/3$: for instance, when Yg appears, Pr will be triggered by Yg to join with probability of about $1/3$.

IV. NOTES ON MODEL EVALUATION

To check the performance of each of our models, we first compare statistics computed with the model to those computed on out-of-sample data. The results for a single choice of in-sample data are shown in Fig. 6. Half of the fights are randomly chosen as in-sample data, with the remaining treated as out-of-sample data to be predicted. We see that the independent model does not capture second- or third-order statistics nor the distribution of fight sizes, while both the equilibrium and dynamic models capture these to produce predictions that are roughly as accurate as using out-of-sample data.

Second, we can check that residuals lie within the bounds of expected statistical fluctu-

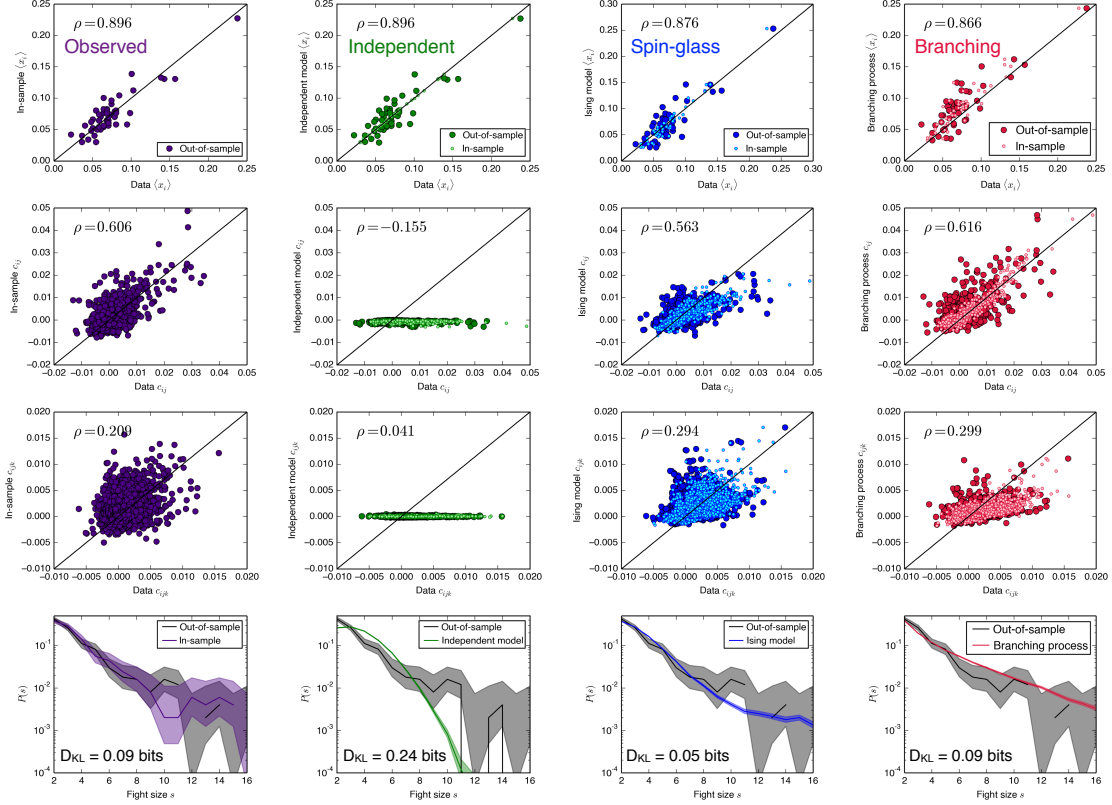


FIG. 6. The degree of fit for the noninteracting model (green), maximum entropy pairwise model (blue), and branching process model (red) to out-of-sample data, compared to the same for the in-sample data (indigo) to which the models are fit. For each model, 10^5 samples were taken to evaluate predicted statistics. Also shown on each plot is the Pearson correlation ρ between predicted and out-of-sample statistics (for individual, pairwise, and triplet statistics) or the Kullback-Leibler divergence D_{KL} between predicted and out-of-sample distributions (for fight sizes). (To avoid problems with large fight sizes that are never observed, D_{KL} is calculated only using fights of size ≤ 12 .)

ations from finite sampling. Shown in Table I, the equilibrium and dynamic models have squared residuals that are below but near the expected value $\langle \chi \rangle^2 = 1$, whereas the independent model is inadequate to describe the statistics. This is visualized in more detail with the distribution of residuals in Fig. 7.

We find no evidence of significant higher order correlations in the data (Fig. 7) and we therefore do not explore models with interactions of higher order. We note however that the resolution of higher order correlations is limited by the finite number of observed fights

TABLE I. Goodness of fit to data for the three models, calculated using Eq. (14) for the independent and pairwise maximum entropy models and Eq. (21) for the dynamic branching model. With $\langle\chi^2\rangle \sim 1$, the equilibrium pairwise and dynamic branching models fit the data roughly within the precision afforded by the data. Overfitting, which would be indicated by $\langle\chi^2\rangle \ll 1$, is avoided by using constrained minimization in the case of the spin-glass model (see Sec. III B) and by ending minimization once $\langle\chi^2\rangle \leq 1$ in the case of the branching model (see Sec. III C).

	<i>Independent model</i>	<i>Equilibrium pairwise model</i>	<i>Dynamic branching model</i>
Random half of data	$\langle\chi^2\rangle = 1.134$	$\langle\chi^2\rangle = 0.459$	$\langle\chi^2\rangle = 0.414$
All data	$\langle\chi^2\rangle = 1.822$	$\langle\chi^2\rangle = 0.597$	$\langle\chi^2\rangle = 0.462$

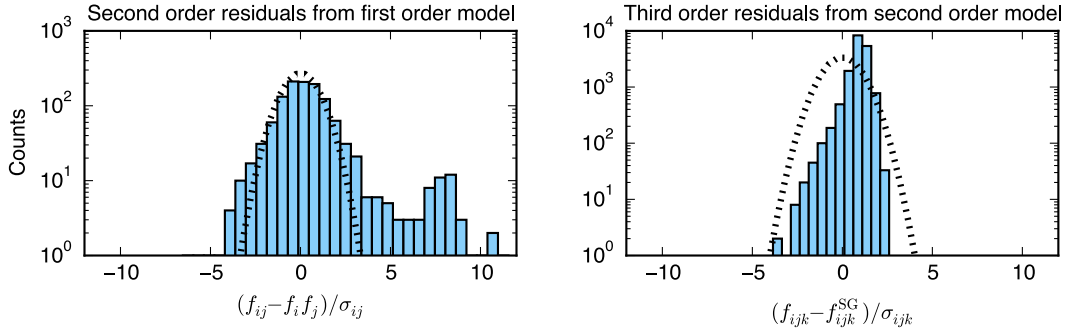


FIG. 7. Second- and third-order statistics in the conflict data. Second-order statistics (left) clearly violate the null expectation for a first-order independent model (dotted line), while third-order statistics (right) lie within expected fluctuations from a second-order model (dotted line). f_{ijk} is the empirical frequency with which each triplet appears in fights, f_{ijk}^{SG} is this frequency in the pairwise equilibrium model, and $\sigma_{ijk} = \sqrt{f_{ijk}(1 - f_{ijk})/N}$ is the expected standard deviation.

and relatively small frequency of individual participation. This cannot easily be remedied by collecting more data as the system is not at equilibrium over longer timescales. To deal with this, we must restrict the data we use in the analyses to collection windows defined by “socially stable periods” (SI, Sec I).

V. NOTES ON EVALUATING SENSITIVITY AND STABILITY

Phase transitions are typically identified as conditions under which the varying of a control parameter causes large-scale changes in the behavior of a system, in a way that sensitivity per individual (measured by, *e.g.*, specific heat or susceptibility) grows arbitrarily large with growing system size. This becomes possible only when there is a collective instability, meaning that the effective size of the perturbation (that starts, say, with a single individual) does not shrink as it spreads through the system but stays of constant size or grows (potentially affecting all individuals). Thus in a finite system, the combination of a peak in sensitivity and collective instability can be used as an indicator of a phase-transition-like state.

A. Sensitivity as Fisher information

In our finite system the notion of diverging sensitivity is arguably more accurately described in terms of information theory. Even when the idea of a phase transition becomes fuzzy in a finite system, the Fisher information measures something adaptively important: the degree to which individual scale perturbations are visible at the global scale, or, equivalently, the connection between the behavior of any individual and the behavior of the whole [S14, S15].

B. Analytical results for sensitivity in the independent model

Here we show that the sensitivity (susceptibility) to increased aggression in the independent model can be efficiently solved numerically. This is used to make a comparison with the pairwise equilibrium model in Fig. 2.

First, in the more analytically straightforward case in which we allow fights of size zero and one ($\alpha = 0$; see section III A), the average fight size and susceptibility are

$$\langle s \rangle_{\alpha=0} = \sum_i (1 + \exp(h_i - h_{\text{ext}}))^{-1} \quad (23)$$

$$\chi_0 = \frac{\partial \langle s \rangle_{\alpha=0}}{\partial h_{\text{ext}}} = \sum_i \text{sech}^2(h_i - h_{\text{ext}}). \quad (24)$$

The partition functions of the constrained and unconstrained models, defined such that

$$p(\vec{x})_{\alpha=0} = \exp[-L_{\alpha=0}(\vec{x})]/Z_0 \quad (25)$$

$$p(\vec{x})_{\alpha \rightarrow \infty} = \exp[-L_{\alpha \rightarrow \infty}(\vec{x})]/Z_\infty, \quad (26)$$

are given by

$$Z_0 = \prod_i (1 + \exp h_i) \quad (27)$$

$$Z_\infty = Z_0 - 1 - \sum_i \exp h_i. \quad (28)$$

In terms of these values, when fights of size zero and one are forbidden ($\alpha \rightarrow \infty$), the average fight size and susceptibility become

$$\langle s \rangle_{\alpha \rightarrow \infty} = \frac{Z_0}{Z_\infty} \langle s \rangle_{\alpha=0} - \frac{\sum_i \exp h_i}{Z_\infty} \quad (29)$$

$$\chi_\infty = \frac{\partial \langle s \rangle_{\alpha \rightarrow \infty}}{\partial h_{\text{ext}}} = \frac{Z_0}{Z_\infty} \frac{\partial \langle s \rangle_{\alpha=0}}{\partial h_{\text{ext}}} - \frac{\sum_i \exp h_i}{Z_\infty} + \frac{Z_0}{Z_\infty} \langle s \rangle_{\alpha=0}^2 - \langle s \rangle_{\alpha \rightarrow \infty}^2. \quad (30)$$

C. Operationalizing collective instability in the branching process model

Collective instability is straightforward to understand in a branching model. In an infinite system, this model has a well-defined phase transition when R_0 , the average number of other individuals triggered when a single individual becomes active, is equal to 1. This also corresponds to the point at which the system is maximally sensitive to changes in parameters. The fact that this “local” amplification factor is also indicative of a global transition relies on the infinite limit: In a finite system, cascades will be shortened when they reach individuals that have already been activated, and maximal sensitivity will happen at some $R_0 > 1$ [S16]. We thus think of R_0 as measuring a “local” or lowest-order stability. In the heterogeneous branching process model, this linear stability of the peaceful state is indicated by the largest eigenvalue R_0 of the redirection probability matrix p_{ij} .

We note that in the branching model, as opposed to the equilibrium model, increasing activation will always eventually lead to instability.

D. Operationalizing collective instability in the pairwise equilibrium model

In an infinite system, the pairwise equilibrium model also has a phase transition under the condition of local instability, with a corresponding diverging sensitivity. In this case, in-

stability can be quantified using the mean-field solution, connecting with a high-temperature expansion of spin-glass models.

One way to think about the continuous phase transition in an infinite spin-glass model is that it is the point at which the high temperature mean-field solution becomes unstable. Mean-field solutions are characterized by frequencies \vec{f} of individual appearance ($f_i = \langle x_i \rangle$) that satisfy the self-consistency equation [S17]

$$f_i = F_i(\vec{f}) \equiv \left[1 + \exp \left(J_{ii} + 2 \sum_{j \neq i} J_{ij} f_j \right) \right]^{-1}. \quad (31)$$

Intuitively, individual i 's frequency of fighting is determined by the mean “field” it feels as a result of others’ fighting. The function F_i encodes how i reacts to its environment, translating the mean fighting frequencies i sees into its own mean frequency. When Eq. (31) holds for every individual using a single set of frequencies \vec{f} , this defines the mean field solution.

Now imagine perturbing fighting frequencies \vec{f} by a small $\Delta\vec{f}$. This will typically no longer be a solution of Eq. (31). But if we repeatedly apply the function F to $\vec{f} + \Delta\vec{f}$, we can imagine two possibilities: we might end up back at \vec{f} (so that $\lim_{n \rightarrow \infty} F^n(\vec{f} + \Delta\vec{f}) = \vec{f}$), or we might get further and further from \vec{f} . We will call the first case a “stable” mean field solution and the second case “unstable”.

For small perturbations $\Delta\vec{f}$, we can distinguish these two cases by taking a derivative to perform a linear stability analysis. Specifically,

$$F_i(\vec{f} + \Delta\vec{f}) \approx F_i(\vec{f}) + \sum_j \frac{\partial F_i}{\partial f_j} \Delta f_j, \quad (32)$$

and to converge back to \vec{f} for every perturbation, we must have that the updated perturbation along each direction has shrunk. This corresponds to a condition on the eigenvalues λ_α of the derivative matrix $M_{ij} \equiv \partial F_i / \partial f_j$; the state is stable if

$$|\lambda_\alpha| < 1 \quad \forall \alpha. \quad (33)$$

Thus the eigenvalue λ with largest magnitude determines stability.

We can write this derivative matrix M more explicitly by taking the derivative of Eq. (31)

and assuming that we are at the fixed point ($\vec{f} = F(\vec{f})$):

$$\begin{aligned}
M_{ij} &= \frac{\partial F_i}{\partial f_j} \\
&= -2(1 - \delta_{ij})J_{ij} \exp\left(h_i + 2 \sum_{j \neq i} J_{ij} f_j\right) f_i^2 \\
&= -2(1 - \delta_{ij})J_{ij} \frac{1 - f_i}{f_i} f_i^2 \\
&= -2(1 - \delta_{ij})J_{ij} f_i (1 - f_i).
\end{aligned} \tag{34}$$

Thus M is a matrix analogous to p_{ij} in the branching process model in that its spectrum is informative about how perturbations grow or shrink. Specifically, we use the magnitude λ of the largest eigenvalue of M as a measure of stability of the system. When $\lambda > 1$, we expect the system to be unstable to perturbation.

This condition on the stability of mean field theory can be shown to be equivalent to the condition that identifies the spin-glass transition in an infinite system. Specifically, instability of the high-temperature mean-field expansion happens only below the spin-glass temperature [S18, S19]. To lowest order in $1/T$ (corresponding to lowest order in J_{ij} or $1/N$), the mean-field free energy has the form (following [S20])

$$A = \sum_i f_i \log f_i + (1 - f_i) \log(1 - f_i) - \sum_i \sum_{j \neq i} J_{ij} f_i f_j - \sum_i J_{ii} f_i, \tag{35}$$

which when differentiated produces the self-consistency equation (31). Taking a second derivative,

$$\frac{\partial^2 A}{\partial m_i \partial m_j} = \delta_{ij} \frac{1}{f_i(1 - f_i)} - (1 - \delta_{ij}) 2J_{ij} \equiv \Lambda_{ij}, \tag{36}$$

which defines stability to this order when all eigenvalues of Λ_{ij} are positive [S19]. Because $f_i(1 - f_i) > 0$, this condition is the same as all eigenvalues of $f_i(1 - f_i)\Lambda_{ij} = \delta_{ij} - M_{ij}$ being positive (where M is defined in Eq. (34)), which is in turn equivalent to the above condition that all eigenvalues of M_{ij} are less than 1.

To create a homogeneous finite system poised at the transition defined by the eigenvalue λ (shown as the red dotted curve in Fig. 2), we define a homogeneous positive h and negative J that make each individual's frequency $f = 1/2$ and $\lambda = 1$.

VI. NOTES ON QUANTIFYING DFC

As described in the main text, by simulating forcing some number of individuals to be involved in every fight we can quantify distance from the critical point by measuring changes to sensitivity and collective instability. We measure this distance in units natural to the system by forcing a subset of individuals to join fights and measuring the sensitivity of the remaining individuals.

The number of individuals (or, in principle, subgroups) that must be forced to reach peak sensitivity provides an operational definition of DFC. We find three to five individuals is sufficient and we find individual variation in sensitivity scores. In both the equilibrium and branching process models individuals with higher sensitivity scores are those who (by definition) exert greater influence on how far the system sits from the critical point.

We find that individuals whose simulated forcing causes the largest increase in average fight sizes are also those who cause the largest increase in sensitivity, as might be expected in this system with largely excitatory interactions. In each model, this sensitivity score depends on the individual's place within the overall network structure, measuring both its direct influence on inducing others to fight and indirect influence through those it induces to fight.

A. Tuning mechanisms

The sensitivity scores generated by each model are capturing different underlying tuning mechanisms. In the case of the branching process model, the model captures the spread of fight-joining through direct interactions: one individual through its behavior triggers the involvement of a second individual. Hence this model only allows that the target individuals, through their own fight joining decisions, can up or down regulate their fight involvement. When a high sensitivity individual up or down regulates its own behavior moving the system respectively closer to or further from the critical point, we call this *direct tuning*.

The equilibrium model, on the other hand, is agnostic to cause, recording any pairwise correlation in fight joining regardless of which individual triggered the joining event. As such it captures the full space of behavioral mechanisms leading to fight-joining—individuals can become involved in fights by up regulating their own fight involvement and *also* through

changes to third-party behavior. For example, policing (by third-parties to the fight) and other conflict management mechanisms reduce the frequency of redirected and directed aggression in the system, and as such can dampen the fight joining behavior of the target individual. When a third party, like a policer, up or down regulates the behavior of high sensitivity individuals we call this *indirect tuning*.

B. Tune towards robustness or towards criticality?

A natural question raised by the discoveries that DFC can be tuned and that there are behavioral mechanisms in place that in principle allow this tuning is whether to tune towards robustness (increase DFC) or criticality (decrease DFC). This decision depends on three factors: (1) the adaptive utility of criticality—when does it make sense to sit near the critical point?, (2) the true state of the environment, and (3) the perceived state of the environment—the accuracy with which system components can detect and encode the environmental state.

The consequences of being near or at the critical point is that information can propagate quickly, with small changes to component behavior inducing large-scale changes in both structure and function. Hence criticality allows the system to more easily reconfigure. How does this work? Moving towards criticality changes the distribution of fight sizes such that there are more large fights. Reconfiguration becomes more likely with large fights because large fights cost more[S10] and costly conflict can lead to changes in alliances, coalitions and the power structure, which controls the cost of conflict management[S3, S8].

We predict that reconfiguration is adaptive when the environment, after having been stable for some time, becomes uncertain. Hence when the state of the environment is stable or very slowly changing and represented by a delta function, the optimal choice (we propose) is to increase DFC. When the environment is uncertain—that is, when it is not clear what the distribution of environmental states is, the system should decrease DFC.

For tuning to be adaptive, the state of the environment must also be correctly perceived by the tuning agent(s). To make this clear, consider the following: a flock of birds in search of food with high and low sensitivity individuals. Some individuals in the flock have poor eyesight and so misjudge the location of food and others have good eyesight. If the high sensitivity birds are also those with poor eyesight, the system may be inappropriately

driven away from food sources. If on the other hand the good eyesight birds are also highly sensitive, the system will be appropriately driven towards food. In a similar way, tuning DFC will require accurate perception of the most beneficial change. *Hence in order for DFC to be tunable and for that tuning to be adaptive in a biological system these two types of heterogeneity must be aligned: high sensitivity score individuals must also be good detectors of environmental state.* Neither of these types of heterogeneity has received much attention in the collective behavior, criticality, or biological phase transition literatures. Consequently little is known about how common it is for these types of heterogeneity to be aligned or whether there are in some systems mechanisms to bring them into alignment.

In our model system, both types of heterogeneity appear to be aligned. Results of earlier work[S3, S5], suggest the individuals with the highest sensitivity scores also are those whose “health state” signals the state of the larger social environment. The idea is as follows: the “health state” of individuals (healthy or compromised) builds up over a history of affiliative and agonistic social interactions with other group members and hence is a slow variable[S21]. The health state of the high sensitivity individuals serves as a proxy for system state because these individuals are among the weakest in the system (lowest 10 percent of the power distribution) and hence register stressful periods more visibly than the other group members [S22].

When these animals are healthy, the system can be said to be in a stable, low variance period. When these individuals are stressed above some baseline level, social dynamics are becoming uncertain. In either case, when aggression during a fight “reaches them”, they move the system closer to the critical point *because of their position in the aggression network as represented by the spin glass and branching process models.* Aggression is less likely to reach them when the policing mechanism is functional and more likely to reach them when the policing mechanism has been disabled [S3].

In Fig. 3, the magenta and blue lines demonstrate the potential heterogeneity of responses when different individuals are forced. The magenta lines demonstrate the effect of forcing individuals in an order that maximizes the resulting average fight size at each step, and blue in an order that minimizes it. The individuals are re-sorted each time another is forced as this can affect the order (for instance, forcing one individual in a strongly-correlated clique can decrease the effect of forcing other individuals in that clique). In Fig. 8, we contrast the case in which individuals are sorted by their effect on the original, unperturbed state of the

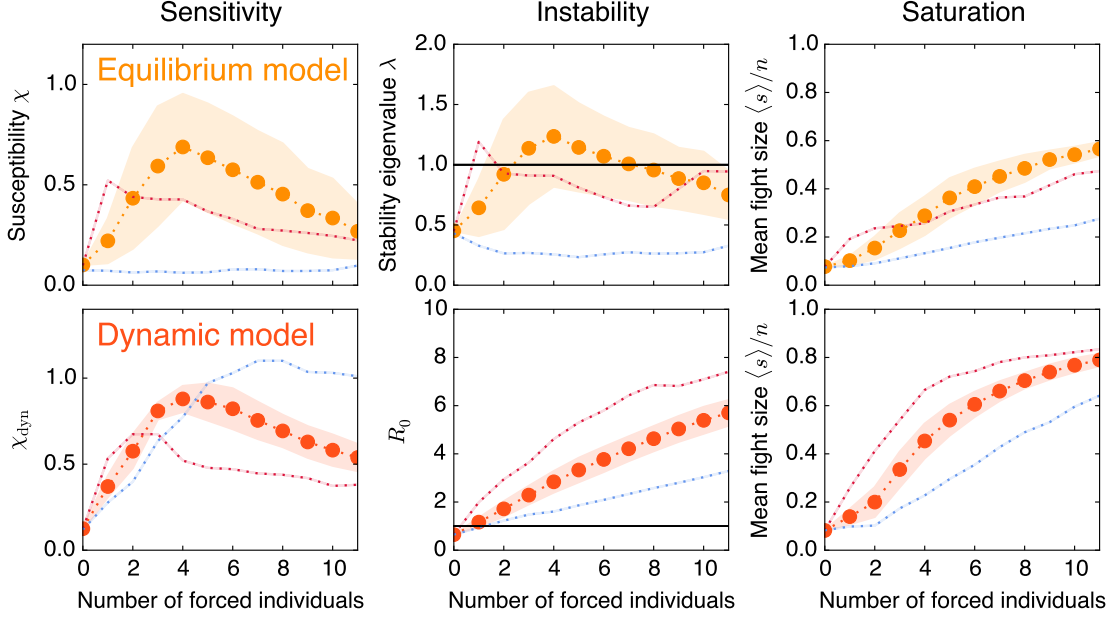


FIG. 8. Same as Fig. 3, except that the magenta and blue lines here show results when forced individuals are chosen as those with the largest and smallest effect on the mean fight size of remaining individuals when forced in the otherwise unperturbed system. (Fig. 3 reperforms this optimization after adding each individual.)

system. The qualitative results are the same as in Fig. 3.

VII. THERMODYNAMIC DERIVATIVES AND FISHER INFORMATION IN THE EQUILIBRIUM MODEL

Quite generally for equilibrium models, it can be shown that the Fisher information is deeply related to important thermodynamic derivatives. The Fisher information is defined as [S23]

$$\mathcal{I}(\mu) = \int \left(\frac{\partial \log p(x)}{\partial \mu} \right)^2 p(x) dx, \quad (37)$$

where μ parameterizes a distribution $p(x)$ describing the behavior of a system, and x represents any number of relevant measurable system state variables. Generalized to multiple parameters, the Fisher information matrix is a fundamental object in information geometry, and forms a Riemannian metric that becomes singular precisely at phase transitions [S15, S24]. $\mathcal{I}(\mu)$ is typically used to measure the amount of information about μ that can

be inferred from draws from $p(x)$. Conversely, if we view individuals as controlling local parameters μ , $\mathcal{I}(\mu)$ measures the degree of control individuals have on group behavior. Then phase transitions, having diverging \mathcal{I} as $N \rightarrow \infty$, correspond to individuals having arbitrarily large effects. But even at finite N , \mathcal{I} measures the amplification of individual information to the global scale. In this sense, \mathcal{I} becomes a straightforward, useful measure of the degree to which a system's behavior is "collective."

Another intuitive meaning comes in terms of the Kullback-Leibler divergence: $\mathcal{I}(\mu)$ represents how quickly the KL-divergence increases as μ is changed, such that

$$D_{\text{KL}}(p(x|\mu)||p(x|\mu + \Delta\mu)) = \mathcal{I}(\mu)(\Delta\mu)^2/2 + O(\Delta\mu^4). \quad (38)$$

Thus the Fisher information measures how quickly the modified distribution becomes distinguishable from the original as μ is varied, and if logs are taken with base 2, $\mathcal{I}(\mu)$ has units of bits per [unit of μ]².

In the case of an equilibrium system described by a Boltzmann distribution, the Fisher information with respect to a local field μ is particularly simple, equal to the derivative of the mean of its conjugate variable x_μ , the generalized susceptibility $\mathcal{I}(\mu) = \frac{\partial}{\partial\mu}\langle x_\mu \rangle$. This example provides a clear link between thermodynamics and information theory. (Yet the Fisher information measure is not limited to equilibrium models, generalizing to dynamic out-of-equilibrium systems by simply interpreting $p(x)$ in Eq. (37) as a distribution over relevant output measurements given some known initial conditions.)

This connection between Fisher information and thermodynamic derivatives is well-established [S15]. Assume we have a system whose distribution over possible states x takes the form of a Boltzmann distribution:

$$p(x) = Z^{-1}e^{-L(x)}. \quad (39)$$

Taking a derivative of $\log p(x)$,

$$\frac{\partial \log p(x)}{\partial \mu} = -\frac{\partial L(x)}{\partial \mu} + Z^{-1} \sum_x \frac{\partial L(x)}{\partial \mu} e^{-L(x)} \quad (40)$$

$$= \left\langle \frac{\partial L}{\partial \mu} \right\rangle - \frac{\partial L(x)}{\partial \mu}, \quad (41)$$

which, when inserted in Eq. (37), gives

$$\mathcal{I}(\mu) = \left\langle \left(\frac{\partial L}{\partial \mu} \right)^2 \right\rangle - \left\langle \frac{\partial L}{\partial \mu} \right\rangle^2. \quad (42)$$

This shows that Fisher information is equal to the variance of the derivative of L . We can further relate this to thermodynamic derivatives by noting that L is typically linearly dependent on certain “fields” (*e.g.* pressure or magnetic field), with derivatives that correspond to measurable macroscopic properties (*e.g.* volume or magnetization). This linearity allows us to write the Fisher information even more simply: when $\langle \partial^2 L / \partial \mu^2 \rangle = 0$,

$$\mathcal{I}(\mu) = -\frac{\partial}{\partial \mu} \left\langle \frac{\partial L}{\partial \mu} \right\rangle. \quad (43)$$

(To see this, explicitly take the derivative of the expectation value:

$$\frac{\partial}{\partial \mu} \left\langle \frac{\partial L}{\partial \mu} \right\rangle = \frac{\partial}{\partial \mu} \left[Z^{-1} \sum_x \exp(-L(x)) \frac{\partial L(x)}{\partial \mu} \right] = - \left\langle \left(\frac{\partial L}{\partial \mu} \right)^2 \right\rangle + \left\langle \frac{\partial L}{\partial \mu} \right\rangle^2 + \left\langle \frac{\partial^2 L}{\partial \mu^2} \right\rangle, \quad (44)$$

which is equal to $-\mathcal{I}(\mu)$ from Eq. (42) when the last term is zero.)

Connecting this result to our equilibrium model, the susceptibility and specific heat are related to the Fisher information with respect to external field h_{ext} and temperature T :

$$\mathcal{I}(h_{\text{ext}}) = \frac{\partial \langle s \rangle}{\partial h_{\text{ext}}} = \chi. \quad (45)$$

$$\mathcal{I}(1/T) = -\frac{1}{T^2} \frac{\partial \langle E \rangle}{\partial T} = C_s. \quad (46)$$

The amount of change in the entire distribution over fights is expressed in terms of a single “order parameter,” the average fight size in the case of varying h_{ext} (and the average energy in the case of varying $1/T$). This implies that if one is trying to infer small changes in the external field h_{ext} by watching the composition of fights, one loses nothing by simply recording the fight sizes.

-
- [S1] J. O. Caldecott. *An Ecological and Behavioural Study of the Pig-Tailed Macaque*. S. Karger, 1986.
- [S2] Jessica C Flack and Frans de Waal. Context modulates signal meaning in primate communication. *Proceedings of the National Academy of Sciences of the United States of America*, 104(5):1581–1586, 2007.
- [S3] J C Flack and D C Krakauer. Encoding power in communication networks. *The American Naturalist*, 168(3):E87–102, September 2006.

- [S4] Jessica C. Flack, M. Girvan, B.M. de Waal, Frans, and David C. Krakauer. Policing stabilizes construction of social niches in primates. *Nature*, 439:426–429, 2006.
- [S5] Jessica C Flack, David C Krakauer, and Frans B M de Waal. Robustness mechanisms in primate societies: a perturbation study. *Proceedings of the Royal Society, Series B Biological sciences*, 272(1568):1091–1099, 2005.
- [S6] Jessica C. Flack, Frans B.M. de Waal, and David C. Krakauer. Social structure, robustness, and policing cost in a cognitively sophisticated species. *American Naturalist*, 165:E126–39, 2005.
- [S7] B Thierry and M Singh. *Macaque Societies: A Model for the Study of Social Organization*. Cambridge University Press, Cambridge, 2004.
- [S8] R. Brush, Eleanor, David C. Krakauer, and Jessica C. Flack. A family of algorithms for encoding consensus about node state in network data. *PLoS Computational Biology*, 9:e1003109, 2013.
- [S9] DC Krakauer, K Page, and JC Flack. The Immuno-Dynamics of Conflict Intervention in Social Systems. *PloS One*, 6:e22709, 2011.
- [S10] Simon DeDeo, David C. Krakauer, and Jessica C. Flack. Inductive game theory and the dynamics of animal conflict. *PLoS Computational Biology*, 6:e1000782, 2010.
- [S11] Elad Schneidman, Michael J Berry II, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007, 2006.
- [S12] John Barton and Simona Cocco. Ising models for neural activity inferred via selective cluster expansion: structural and coding properties. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(03):P03002, March 2013.
- [S13] Simon DeDeo, David C Krakauer, and Jessica C Flack. Inductive Game Theory and the Dynamics of Animal Conflict. *PLoS Comp. Bio.*, 6:e1000782, 2010.
- [S14] Martin Tchernookov and Ilya Nemenman. Predictive information in a nonequilibrium critical model. pages 1–16, 2012.
- [S15] Mikhail Prokopenko, Joseph T. Lizier, Oliver Obst, and X. Rosalind Wang. Relating Fisher information to order parameters. *Physical Review E*, 84(4):041116, October 2011.
- [S16] John M Beggs. The criticality hypothesis: how local cortical networks might optimize information processing. *Philosophical transactions. Series A, Mathematical, physical, and engi-*

- neering sciences*, 366(1864):329–43, February 2008.
- [S17] H E Stanley. *Introduction to phase transitions and critical phenomena*. Oxford University Press, 1971.
 - [S18] M Mezard, G Parisi, and M Virasoro. *Spin Glass Theory And Beyond*, volume 9 of *World Scientific Lecture Notes in Physics*. World Scientific, 1987.
 - [S19] Antoine Georges. Strongly Correlated Electron Materials: Dynamical Mean-Field Theory and Electronic Structure. In A Avella and F Mancini, editors, *Lectures on the Physics of Highly Correlated Electron Systems VIII: Eighth Training Course*, volume 3, page 71. 2004.
 - [S20] Antoine Georges and Jonathan S Yedidia. How to expand around mean-field theory using high-temperature expansions. *Journal of Physics A: Mathematical and General*, 24:2173–2192, 1991.
 - [S21] Jessica C. Flack, Doug Erwin, Tanya Elliot, and David C. Krakauer. Timescales, symmetry, and uncertainty reduction in the origins of hierarchy in biological systems. *Cooperation and its Evolution*, pages 45–74, 2013.
 - [S22] Jessica C. Flack, P. Hammerstein, and D.C. Krakauer. Robustness in biological and social systems. In P. Hammerstein and J.R. Stevens, editors, *Evolution and the Mechanisms of Decision-making*. MIT Press, 2012.
 - [S23] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 1991.
 - [S24] Gavin Crooks. Measuring Thermodynamic Length. *Physical Review Letters*, 99(10):100602, September 2007.